

IN THE SPECIFICATION:

Please amend the paragraph beginning at page 8, line 7, as follows.

5 The allocation of various domains to each proxy server 120 attempts to assign the heavy file types and the heavy domains to individual proxy servers 120 with the aim of separating out the larger requests as well as balancing the overall load. If there are P proxies and the heavy file types account for fraction $1/h$ of the total load, then we assign $P \times 1/h$ of the proxy servers 120 to exclusively serve heavy file types. The heavy domains are sorted in increasing order of their average file sizes; we then split this list into $P \times (1 - (1/h))$ partitions of equal load, and assign one partition to each of the remaining proxy servers 120. Here, the load for heavy domains is computed after excluding the requests that are of heavy types. We assume that all proxy servers 120 have identical capacities; otherwise, the load can be spread in proportion to their capacities and the scheme works with no significant variation. The motivation for separating heavy types and sorting heavy domains by size is to reduce the variance in request sizes arriving at each proxy server 120-i, 15 since large variances can affect the slowdown of tasks in the request queue. The effect of task size variance on the slowdown depends on the scheduling policy at the request queue. For example, with a FCFS policy, slowdown is proportional to variance.